arXiv:1607.06139v1 [cs.DC] 20 Jul 2016

A Complexity-Based Hierarchy for Multiprocessor Synchronization

Faith Ellen
University of Toronto
faith@cs.toronto.edu

Rati Gelashvili
MIT
gelash@mit.edu

Nir Shavit
MIT
shanir@csail.mit.edu

Leqi Zhu
University of Toronto
lezhu@cs.toronto.edu

Abstract

For many years, Herlihy's elegant computability based Consensus Hierarchy has been our best explanation of the relative power of various types of multiprocessor synchronization objects when used in deterministic algorithms. However, key to this hierarchy is treating synchronization instructions as distinct objects, an approach that is far from the real-world, where multiprocessor programs apply synchronization instructions to collections of arbitrary memory locations. We were surprised to realize that, when considering instructions applied to memory locations, the computability based hierarchy collapses. This leaves open the question of how to better capture the power of various synchronization instructions.

In this paper, we provide an approach to answering this question. We present a hierarchy of synchronization instructions, classified by their space complexity in solving obstruction-free consensus. Our hierarchy provides a classification of combinations of known instructions that seems to fit with our intuition of how useful some are in practice, while questioning the effectiveness of others. We prove an essentially tight characterization of the power of buffered read and write instructions. Interestingly, we show a similar result for multi-location atomic assignments.

1 Introduction

Herlihy's Consensus Hierarchy [Her91] assigns a consensus number to each object, namely, the number of processes for which there is a wait-free binary consensus algorithm using only instances of this object and read-write registers. It is simple, elegant and, for many years, has been our best explanation of synchronization power.

Robustness says that, using combinations of objects with consensus numbers at most k, it is not possible to solve wait-free consensus for more than k processes [Jay93]. The implication is that modern machines need to provide objects with infinite consensus number. Otherwise, they will not be universal, that is, they cannot be used to implement all objects or solve all tasks in a wait-free (or non-blocking) manner for any number of processes [Her91, Tau06, Ray12, HS12]. Although there are ingenious non-deterministic constructions that prove that Herlihy's Consensus Hierarchy is not robust [Sch97, LH00], it is known to be robust for deterministic one-shot objects [HR00] and deterministic read-modify-write and readable objects [Rup00]. It is unknown whether it is robust for general deterministic objects.

In adopting this explanation of computational power, we failed to notice an important fact: multiprocessors do not compute using synchronization objects. Rather, they apply synchronization instructions to locations in memory. With this point of view, Herlihy's Consensus Hierarchy no longer captures the phenomena we are trying to explain.

For example, consider two simple instructions:

- *fetch-and-add*(2), which returns the number stored in a memory location and increases its value by 2, and

- *test-and-set*(), which returns the number stored in a memory location and sets it to 1 if it contained 0[1].

Objects that support only one of these instructions as an operation have consensus number 2 and cannot be combined to solve wait-free consensus for 3 or more processes. However, a system that supports both instructions can solve wait-free binary consensus for any number of processes. The protocol uses a single memory location initialized to 0. Processes with input 0 perform *fetch-and-add*(2), while processes with input 1 perform *test-and-set*(). If the value returned is odd, the process decides 1. If the value 0 was returned from *test-and-set*(), the process also decides 1. Otherwise, the process decides 0.

Another example considers three instructions:

- *read*(), which returns the number stored in a memory location,
- *decrement*(), which decrements the number stored in a memory location and returns nothing, and
- *multiply*(x), which multiplies the number stored in a memory location by x and returns nothing.

A similar situation arises: Objects that support *read* and one of *decrement* and *multiply* have consensus number 1 and cannot be combined to solve wait-free consensus for 2 or more processes. However, a system that supports all three instructions can solve wait-free binary consensus for any number of processes. The protocol uses a single memory location initialized to 1. Processes with input 0 perform *decrement*(), while processes with input 1 perform *multiply*(n). The second operation by each process is *read*(). If the value returned is positive, then the process decides 1. If it is negative, then the process decides 0.

For randomized computation, Herlihy's Consensus Hierarchy also collapses: randomized wait-free binary consensus among any number of processes can be solved using only read-write registers, which have consensus number 1. Fich, Herlihy, and Shavit [FHS98] proved that $\Omega(\sqrt{n})$ historyless objects, which support only trivial operations, such as *read*, and historyless operations, such as *write*, *test-and-set*, and *swap*, are necessary to solve this problem. They noted that, in contrast, only one fetch-and-increment or fetch-and-add object suffices for solving this problem. Yet, these objects and historyless objects are similarly classified in Herlihy's Consensus Hierarchy (i.e. they all have consensus number 1 or 2). They suggested that the number of instances of an object needed to solve randomized wait-free consensus among n processes might be another way to classifying its power.

Motivated by these observations, we consider a classification of instruction sets based on the number of memory locations needed to solve *obstruction-free n-valued* consensus among $n \geq 2$ processes. Obstruction freedom is a simple and natural progress measure. Some state-of-the-art synchronization operations, for example hardware transactions [Int12], do not guarantee more than obstruction freedom. Obstruction freedom is also closely related to randomized computation. In fact, any (deterministic) obstruction free algorithm can be transformed into a randomized wait-free algorithm that uses the same number of memory locations (against an oblivious adversary) [GHHW13]. Obstruction-free algorithms can also be transformed into wait-free algorithms in the unknown-bound semi-synchronous model [FLMS05].

1.1 Our Results

For various instruction sets $\mathcal{I}$, we provide upper and lower bounds on $\mathcal{SP}(\mathcal{I}, n)$, the number of memory locations (supporting $\mathcal{I}$) needed to solve obstruction-free n-valued consensus among $n \geq 2$ processes (abbreviated as *n-consensus*). The results are summarized in Table 1.

We begin, in Section 3, by considering the instructions

- *multiply*(x), which multiplies the number stored in a memory location by x and returns nothing, and
- *add*(x), which adds x to the number stored in a memory location and returns nothing. and
- *set-bit*(x), which sets bit x of a memory location to 1 and returns nothing.

[1]This definition of *test-and-set* is slightly stronger than the standard one, which simply always sets the location to 1. However, they have the same consensus number and they behave identically if the values are always binary.

We show that one memory location supporting *read*() and one of these instructions can be used to solve n-consensus. The idea is to show that these instruction sets can implement n counters in a single location. We can then use a *racing counters* algorithm.

Next, we consider *max-registers* [AAC09], which are memory locations supporting

- *read-max*(), which reads the number stored in a memory location, and
- *write-max*(x), which stores the number x in a memory location, provided it contains a value less than x, and returns nothing.

In Section 4, we prove that two max registers are necessary and sufficient for solving n-consensus.

In Section 5, we prove that a single memory location supporting $\{read(), write(x), fetch\text{-}and\text{-}increment()\}$ cannot be used to solve n-consensus, for $n \geq 3$. We also present an algorithm for solving n-consensus using $O(\log n)$ such memory locations.

In Section 6, we introduce a family of instructions $\mathcal{B}_\ell$, for $\ell \geq 1$, and show how to solve n-consensus using $\lceil \frac{n}{\ell} \rceil$ memory locations supporting these instructions. Extending Zhu's $n-1$ lower bound [Zhu16], we also prove that $\lceil \frac{n-1}{\ell} \rceil$ such memory locations are necessary, which is tight except when $n-1$ is divisible by ℓ.

Our main technical contribution is in Section 7, where we show that the preceding lower bound holds within a factor of two even in the presence of atomic multiple assignment. Multiple assignment can be implemented by simple transactions, so our result implies that such transactions cannot significantly reduce space complexity. The proof further extends the techniques of [Zhu16] via a nice combinatorial argument, which we hope will be of independent interest.

There are algorithms that solve n-consensus using n registers [AH90, BRS15, Zhu15]. The best-known lower bound is $n-1$ registers [Zhu16]. A modification of the anonymous algorithm for n-consensus in [Zhu15] allows us to use only $n-1$ memory locations supporting $\{read(), swap(x)\}$. It is presented in Section 8. We also explain how to adapt the lower bound in [Zhu16] to get a simpler proof that $\Omega(\sqrt{n})$ memory locations, which support only trivial and historyless instructions, are necessary to solve n-consensus.

Finally, in Section 9, we show that an unbounded number of memory locations supporting *read*() and either *write*(*1*) or *test-and-set*() are necessary and sufficient to solve n-consensus, for $n \geq 3$. Furthermore, we show how to reduce the number of memory locations to $O(n \log n)$ when *write*(0) is also available.

Instructions $\mathcal{I}$	$\mathcal{SP}(\mathcal{I}, n)$
$\{read(), test\text{-}and\text{-}set()\}$, $\{read(), write(1)\}$	∞
$\{read(), test\text{-}and\text{-}set(), reset()\}$, $\{read(), write(1), write(0)\}$	$n-1$ (lower), $O(n \log n)$ (upper)
$\{read(), write(x)\}$	$n-1$ (lower), n (upper)
$\{read(), swap(x)\}$	$\Omega(\sqrt{n})$ (lower), $n-1$ (upper)
$\{\ell\text{-}buffer\text{-}read(), \ell\text{-}buffer\text{-}write(x)\}$	$\lceil \frac{n-1}{\ell} \rceil$ (lower), $\lceil \frac{n}{\ell} \rceil$ (upper)
$\{read(), write(x), increment()\}$ $\{read(), write(x), fetch\text{-}and\text{-}increment()\}$	2 (lower), $O(\log n)$ (upper)
$\{read\text{-}max(), write\text{-}max(x)\}$	2
$\{compare\text{-}and\text{-}swap(x, y)\}$ $\{read(), set\text{-}bit(x)\}$ $\{read(), add(x)\}$, $\{read(), multiply(x)\}$ $\{fetch\text{-}and\text{-}add(x)\}\}$, $\{fetch\text{-}and\text{-}multiply(x)\}$	1

Table 1: Space Hierarachy

2 Model

Our model bears similarities to the standard asynchronous shared memory model, albeit with a few important differences. We consider a system of $n \geq 2$ processes that *supports* some set of deterministic synchronization *instructions*, $\mathcal{I}$, on a set of identical memory locations. The processes take steps at arbitrary, possibly changing, speeds and may crash at any time. Each step is an atomic invocation of some instruction on some memory location. Scheduling is controlled by an adversary.

The processes can use instructions on the memory locations to *simulate* (or *implement*) various objects. An object provides a set of operations which processes can call. Although a memory location together with the supported instructions can be viewed as an object, we do not do so, to emphasize the *uniformity requirement* that the same set of instructions is supported on all memory locations.

We consider the problem of solving *obstruction-free m-valued consensus* in such a system. Here, each of the n processes has an input from $\{0, 1, \ldots, m-1\}$ and is supposed to output the same input value (called a *decision*). Obstruction-free means that every process can eventually decide a value if no other process is taking steps at the same time. When $m = n$, we will refer to this problem as *n-consensus*.

In some of our algorithms, we let the processes take snapshots of all memory locations. If the system supports *read*(), or a similar instruction, then processes can implement an atomic snapshot in an obstruction-free way by using the double collect algorithm in [AAD+93]: i.e. they repeatedly collect the values of all locations until they observe two adjacent collects whose values agree in each location.

In our lower bound proofs, we consider solving n-consensus when each process starts with an input in $\{0, 1\}$ (i.e. binary consensus). We say that a process p is *poised* to perform a specific instruction on a memory location r at a configuration C if p's next step in configuration C is to perform that instruction to r. Since we consider deterministic algorithms, the next step of a process is uniquely defined at every configuration.

3 Arithmetic Instructions

Consider a system that supports only *read*() and either *add*(x) or *multiply*(x). We show how to solve n-consensus using a single memory location in such a system. The idea is to show that we can simulate certain collections of objects that can solve n-consensus.

A *counter* object supports two operations, *read*() and *increment*(), where *read*() returns the current value of the counter and *increment*() increments the value of the counter by 1. In the next lemma, we present a *racing counters* protocol that bears some similarity to the consensus protocol of Aspnes and Herlihy [AH90].

Lemma 3.1. *It is possible to solve obstruction-free m-valued consensus among any number of processes using m counter objects.*

Proof. We associate a separate counter c_v for each possible input value v. Processes that *promote* a particular value v increment c_v. (Initially each process promotes its input value.) After performing an increment, a process takes an atomic snapshot of all m counters and promotes the value associated with a counter with the maximum value (breaking ties arbitrarily). When a process observes that some counter c_v is more than n larger than all other counters, it returns the value v associated with c_v. This works because all other processes will increment some counter at most once before next taking a snapshot, where c_v will still be the only maximum. From there on, all processes will promote value v and keep incrementing c_v, which will shortly become large enough for all processes to return v. Obstruction-freedom follows because a process running on its own will continue to increment the same counter and, thus, eventually be able to return. □

The counters may grow arbitrarily large in the preceding protocol. The next lemma shows that it is possible to overcome this limitation when the counter objects additionally support a *decrement*() operation, which simply decrements the value of the counter object by 1. We call such counter objects *decrementable*.

Lemma 3.2. *It is possible to solve obstruction-free m-valued consensus among any number of processes using m decrementable counter objects, each of which only takes on values in the range $\{0, 1, \ldots, 3n-1\}$.*

Proof. We modify the construction in Lemma 3.1 slightly by changing what a process does when it wants to increment c_u after performing a snapshot to promote the value u. Let c_v be a counter with the largest value among the other counters (i.e. excluding c_u, which has the largest value). If $c_v < n$, it increments c_u, as before. If $c_v \geq n$, then, instead of incrementing c_u, it decrements c_v.

A counter with value 0 is never decremented. This is because, after the last time some process observed it to have value at least n, each process will decrement the counter at most once before reading its value again. Similarly, a counter c_u never becomes larger than $3n-1$. After the last time some process observed it to have value less than $2n$, each process can increment c_u at most once before reading its value again. If $c_u \geq 2n$, then either the largest value of the other counters is less than n, in which case the process returns

without incrementing c_u, or it is at least n, in which case the process decrements a counter with this value, instead of incrementing c_u. □

In the following theorem, we show how to n simulate counters and n decrementable counters.

Theorem 3.3. *It is possible to solve n-consensus using a single memory location that supports only read() and either multiply(x) or add(x).*

Proof. We first give an obstruction-free implementation n counters using a single location in a system with *read*() and *multiply*(x) instructions, which proves the claim for this set of instructions, by Lemma 3.1. The location is initialized with 1. For a value v, let p_v be the $(v+1)$'st prime number. A process increments a counter corresponding to v by performing *multiply*(p_v). A *read*() instruction returns the value x currently stored in the memory location, which gives a snapshot of all counter values: The value of the counter c_v is the exponent of p_v in the prime decomposition of x.

A similar construction does not work in a system with *read*() and *add*(x) instructions. For example, suppose one counter is incremented by calling *add*(a) and another counter is incremented by calling *add*(b). Then the value ab can be obtained by incrementing the first counter b times or incrementing the second counter a times. However, we can use a single memory location that supports $\{read(), add(x)\}$ to implement n different decrementable counters which always have values in $\{0, 1, \ldots, 3n-1\}$, which is also sufficient for solving consensus by Lemma 3.2. We initialize the location with 0. We view the value stored in the location as a number written in base $3n$ and interpret its i'th least significant digit as the value of counter c_{i-1}. To increment c_i, we perform *add*($(3n)^i$), to decrement c_i, we perform *add*($-(3n)^i$), and *read*() provides an atomic snapshot of all n counters. □

4 Max-Registers

A *max-register* object [AAC09] supports two operations, *write-max*(x) and *read-max*(). The *write-max*(x) operation sets the value of the max-register to x if x is larger than the current value and *read-max*() returns the current value of the max-register (which is the largest amongst all values previously written to it). We show that 2 max-registers are necessary and sufficient for solving n-consensus.

Theorem 4.1. *It is not possible to solve n-consensus using a single max-register.*

Proof. Consider a solo execution α of the first process that returns 0, and a solo execution β of the second process that returns 1. We show how to interleave these two executions so that the resulting execution is indistinguishable to both processes from their respective solo executions. Hence, both values will be returned, violating the agreement property of consensus.

To build the interleaved execution, run both processes until they are first poised to perform *write-max*. Suppose the first process is poised to perform *write-max*(a) and the second process is poised to perform *write-max*(b). If $a \leq b$, let the first process take steps until it is next poised to perform *write-max* or until the end of α, if it performs no more *write-max* operations. Otherwise, let the second process take steps until it is next poised to perform *write-max* or until the end of β. Repeat this until one of the processes reaches the end of its execution and then let the other process finish. □

Theorem 4.2. *It is possible to solve m-valued consensus for any number of processes using only two max registers.*

Proof. We describe an obstruction-free protocol for m-valued consensus using two max registers, m_1 and m_2. Consider the lexicographic ordering $\prec$ on the set $S = \{(r, x) : r \geq 0 \text{ and } x \in \{0, \ldots, m-1\}\}$. Let y be a fixed prime that is larger than m. Note that, for $(r, x), (r', x') \in S$, $(r, x) \prec (r', x')$ if and only if $(x+1)y^r < (x'+1)y^{r'}$. Thus, by identifying $(r, x) \in S$ with $(x+1)y^r$, we may assume that m_1 and m_2 are max registers defined on S with respect to the lexicographic ordering $\prec$.

Initially, both m_1 and m_2 are $(0, 0)$. A process begins by performing a write-max to m_1 with $(0, x)$, where $x \in \{0, \ldots, m-1\}$ is its input value. Subsequently, it repeatedly takes snapshots of m_1 and m_2 via an obstruction-free double-collect (using read-max instead of read). If it sees a snapshot where $m_1 = (r+1, x)$ and $m_2 = (r, x)$, then it decides x and terminates. If it sees a snapshot where $m_1 = m_2 = (r, x)$, then it

performs a write-max to m_1 with value $(r+1, x)$. Otherwise, it performs a write-max to m_2 with the value it saw in m_1.

We now show that no two processes can decide different values. Consider the first process p to decide a value, say x. Then p took a snapshot where $m_1 = (r+1, x)$ and $m_2 = (r, x)$, for some $r \geq 0$. Suppose, for a contradiction, that some other process q decides $x' \neq x$. Then q took a snapshot where $m_1 = (r'+1, x')$ and $m_2 = (r', x')$, for some $r' \geq 0$. Since p was the first process to decide a value, q's snapshot occurs after p's snapshot, so $(r, x) \preceq (r', x')$. Since $x' \neq x$, it follows that $(r, x) \prec (r', x')$.

We observe the following key property: Since q took a snapshot where $m_1 = (r'+1, x')$, then for $0 \leq j \leq r'$, there was a process p_j that took a snapshot where $m_1 = m_2 = (j, x)$. Otherwise, consider the smallest j such that $0 \leq j \leq r'$ and no process took a snapshot where $m_1 = m_2 = (j, x')$. By definition of j or by assumption (if $j = r'$), there was a process that took a snapshot where $m_1 = (j+1, x')$. Hence, some process performed a write-max of $(j+1, x')$ to m_1. From the algorithm, this process took a snapshot where $m_1 = m_2 = (j, x')$, which is a contradiction.

Now consider the smallest j such that $0 \leq j \leq r'$ and $(r, x) \prec (j, x')$. Note that $m_2 = (r, x)$ in p's snapshot, while $m_2 = (j, x')$ in p_j's snapshot. Since $(r, x) \prec (j, x')$, it follows that p_j's snapshot occurs after p's snapshot. But $m_1 = (r+1, x)$ in p's snapshot, $m_1 = (j, x')$ in p_j's snapshot, and p_j's snapshot occurs after p's snapshot, so $(r+1, x) \preceq (j, x')$. Since $x \neq x'$, it follows that $(r+1, x) \prec (j, x')$. This means that $j \geq 1$ and $(r, x) \prec (j-1, x')$, which contradicts the fact that j was the smallest. □

5 Increment

Consider a system that supports only *read*, *write*(x), and either *fetch-and-increment*() or *increment*(). We prove that it is not possible to solve n-consensus in the first (stronger) case using a single memory location and we provide an algorithm in the second (weaker) case, which uses $O(\log n)$ memory locations.

Theorem 5.1. *It is not possible to solve n-consensus using a single memory location that supports only read(), write(x), and fetch-and-increment().*

Proof. Suppose there is a binary consensus algorithm for two processes, p and q, using only 1 memory location. Consider solo executions α and β by p with input 0 and input 1, respectively. Let α' be the longest prefix of α that does not contain a *write* and define β' analogously. Without the loss of generality, assume that β' performs at least as many *fetch-and-increment* instructions as α'. Let C be the configuration that results from performing α' starting from the initial configuration in which p has input 0 and the other process, q has input 1.

Consider the shortest prefix β'' of β' in which p performs the same number of fetch-and-increments as it performs in α'. Let C' be the configuration that results from performing β'' starting from the initial configuration in which both p and q have input 1. Then q must decide 1 in its solo execution γ starting from configuration C'. However, C and C' are indistinguishable to process q, so it must decide 1 in γ starting from configuration C. Thus, p cannot have decided yet in configuration C, otherwise both 0 and 1 would have been decided.

Therefore, p is poised to perform a *write* in configuration C. Let α'' be the remainder of α, so $\alpha = \alpha'\alpha''$. The configurations resulting from performing this *write* starting from C and $C\gamma$ are indistinguishable to p, so p also decides 0 starting from $C\gamma$. But in this execution, both 0 and 1 are decided. This is a contradiction. □

The following well-known construction converts any algorithm for solving binary consensus to an algorithm for solving n-valued consensus [HS12].

Lemma 5.2. *Consider a system that supports a set of instructions that includes read() and write(x). If it is possible solve obstruction-free binary consensus among n processes using only c memory locations, then it is possible to solve n-consensus using only $(c+2) \cdot \lceil \log_2 n \rceil - 2$ locations.*

Proof. The processes agree bit-by-bit in $\lceil \log_2 m \rceil$ layers, each consisting of $c+2$ locations. A process starts in the first layer with its input value as its value for layer 1. In layer i, If the i'th bit of its value is 0, a process records this value in the designated 0-location for the layer. Otherwise, it records it in the designated 1-location. Then, it performs the obstruction-free binary consensus algorithm using c locations to agree on

the i'th bit v of the output. If this bit differs from the i'th bit of its value, the process reads one of the recorded values from the designated v-location for layer i and adopts its value for the next round. Note that some process must have already recorded a value to this location since, otherwise, the bit $\bar{v}$ would have been agreed upon. This ensures that the values used for layer $i+1$ are all input values and they all agree in their first i bits. By the end, all processes have agreed on $\lceil \log_2 m \rceil$ bits, i.e. on one of the at most m different input values.

We can save two locations because the last layer does not require designated 0 and 1-locations. □

By Lemma 3.1, n processes can solve obstruction-free binary consensus using two counter objects and, hence, in this system, using two memory locations. The next result then follows from Lemma 5.2.

Theorem 5.3. *It is possible to solve n-consensus using only $O(\log n)$ memory locations that support only read(), write(x), and increment().*

6 Buffers

In this section, we consider the instructions *ℓ-buffer-read()* and *ℓ-buffer-write(x)*, for $\ell \geq 1$, which generalize read and write, respectively. Specifically, an *ℓ-buffer-read()* instruction returns the last ℓ inputs to *ℓ-buffer-write* instructions previously applied on the memory location, in order from least recent to most recent.

We consider a system that supports the instruction set $\mathcal{B}_\ell = \{\ell\text{-buffer-read}(), \ell\text{-buffer-write}(x)\}$, for some $\ell \geq 1$. We call each memory location in such a system an *ℓ-buffer*. Note that a 1-buffer is simply a register. For $\ell > 1$, an ℓ-buffer essentially maintains a buffer of the ℓ most recent writes to that location and allows them to be read.

In Section 6.1, we show that a single ℓ-buffer can be used to simulate a powerful *history* object that can be updated by at most ℓ processes. This will immediately allow us to simulate Aspnes and Herlihy's algorithm for n-consensus [AH90] and, hence, solve n-consensus, using only $\lceil n/\ell \rceil$ ℓ-buffers. In Section 6.2, we prove that $\lceil (n-1)/\ell \rceil$ ℓ-buffers are necessary, which matches the upper bound whenever $n-1$ is not a multiple of ℓ.

6.1 Simulations Using Buffers

A *history* object H supports two operations, *get-history()* and *append(x)*, where *get-history()* returns the sequence of all values appended to H by prior *append* operations, in order. We first show that, using a single ℓ-buffer B, we can simulate a history object H that supports at most ℓ different appenders, but arbitrarily many readers.

Lemma 6.1. *A single ℓ-buffer can simulate a history object on which at most ℓ different processes can perform append(x) and any number of processes can perform get-history().*

Proof. Without loss of generality, assume that no value is appended to H more than once. This can be achieved by having a process include its process identifier and a sequence number along with the value that it wants to append.

In our implementation, B is initially $\perp$ and each value written to B is of the form $(\mathbf{h}, x)$, where $\mathbf{h}$ is a history of appended values and x is a single appended value.

To *get-history()* from H, a process simply performs an *ℓ-buffer-read* of B to obtain a vector $(a_1, \ldots, a_\ell)$. The operation is linearized at this step. If $(a_1, \ldots, a_\ell) = (\perp, \ldots, \perp)$, then the empty sequence is returned. Otherwise, let k be the smallest integer such that $a_k \neq \perp$. It follows that, for $k \leq i \leq \ell$, $a_i = (\mathbf{h}_i, x_i)$, where $\mathbf{h}_i$ is a history and x_i is an appended value. If $k > 1$, then $(x_k, x_{k+1}, \cdots, x_\ell)$ is returned. Now suppose $k = 1$. Let $\mathbf{h}$ be the longest history amongst $\mathbf{h}_1, \ldots, \mathbf{h}_\ell$. If $\mathbf{h}$ contains x_1, then $\mathbf{h}' \cdot (x_1, \ldots, x_\ell)$ is returned, where $\mathbf{h}'$ is the prefix of $\mathbf{h}$ up to, but not including, x_1. If $\mathbf{h}$ does not contain x_1, then $\mathbf{h} \cdot (x_1, \ldots, x_\ell)$ is returned.

To *append(x)* to H, a process first performs a *get-history()* on H to obtain a history $\mathbf{h}$ and then performs *ℓ-buffer-write* on B with value $(\mathbf{h}, x)$. The operation is linearized at this *ℓ-buffer-write* step.

We claim that every *get-history()* operation X returns the sequence of inputs to all *append* operations on H linearized before the linearization point of X, in order from least recent to most recent. The proof is

by induction on the number of *append* operations on H linearized before X. Let R be the ℓ-*buffer-read* step in X, which is the linearization point of X, and let $(a_1, \ldots, a_\ell)$ be the vector returned by R.

If there are no *append* operations linearized before R, then there are no ℓ-*buffer-write* steps before R. It follows that $(a_1, \ldots, a_\ell) = (\perp, \ldots, \perp)$ and X returns the empty sequence, as required.

Now, suppose that there is at least one *append* operation linearized before R. Let k be the smallest integer such that $a_k \neq \perp$. For $k \leq i \leq \ell$, $a_i = (\mathbf{h}_i, x_i)$, where $\mathbf{h}_i$ is a history returned by a *get-history*() operation X_i on H which had strictly fewer *append* operations linearized before it than X.

If $k > 1$, then there are only $\ell - k + 1$ occurrences of ℓ-*buffer-write* before R. Since each *append* operation is linearized at its ℓ-*buffer-write*, only $\ell - k + 1$ *append* instructions are linearized before X and $x_k, \ldots, x_\ell$ are the values appended by these $\ell - k + 1$ *append* operations, in order from least recent to most recent. In this case, $(x_k, x_{k+1}, \ldots, x_\ell)$ is returned, as required.

Now suppose $k = 1$. Let $\mathbf{h}$ be the longest history amongst $\mathbf{h}_1, \ldots, \mathbf{h}_\ell$, let m be such that $\mathbf{h}_m = \mathbf{h}$, and let W be the ℓ-*buffer-write* which wrote x_1 to B. There are two cases to consider.

Case 1: $\mathbf{h}$ *contains* x_1. By the induction hypothesis, X_m returns the history of all values appended to H by *append* operations linearized before X_m, in order from least recent to most recent. By definition, $x_1, \ldots, x_\ell$ are the last ℓ values appended to H prior to R, in order. Since $\mathbf{h}$ contains x_1, it must also contain all values appended to H prior to x_1. It follows that $\mathbf{h}' \cdot (x_1, \ldots, x_\ell)$ is the correct return value, where $\mathbf{h}'$ is the prefix of $\mathbf{h}$ up to, but not including, x_1.

Case 2: $\mathbf{h}_1$ *does not contain* x_1. Since $\mathbf{h}_i$ is a prefix of $\mathbf{h}$, it does not contain x_1, either. By the induction hypothesis, W occurs after $X_1, \ldots, X_\ell$. Thus, the *append* operations which performed $X_1, \ldots, X_\ell$ are all concurrent. In particular, W occurs in each of their execution intervals. It follows that these *append* operations are all by different processes. Since there are at most ℓ updating processes, there is no *append* operation linearized between X_m and W. Therefore, $\mathbf{h}$ contains all values appended to H prior to W. It follows that $\mathbf{h} \cdot (x_1, \ldots, x_\ell)$ is the correct return value. □

In fact, this lemma allows us to simulate *any* object that supports at most ℓ updating processes using only a single ℓ-buffer. This is because the state of an object is determined by the history of the non-trivial operations performed on it. In particular, we can simulate an array of ℓ single-writer registers using a single ℓ-buffer.

Lemma 6.2. *A single ℓ-buffer can simulate ℓ single-writer registers.*

Proof. Suppose that register R_i is owned by process p_i, for $1 \leq i \leq \ell$. By Lemma 6.1, it is possible to simulate a history object H that can be updated by ℓ processes and read by any number of processes. To write value x to R_i, process p_i appends (i, x) to H. To read R_i, a process reads H and finds the value of the most recent write to R_i. This is the second component of the last pair in the history whose first component is i. □

Thus, we can use $\lceil \frac{n}{\ell} \rceil$ ℓ-buffers to simulate n single-writer registers. This, in turn, allows us to simulate the n-consensus algorithm of Aspnes and Herlihy [AH90].

Theorem 6.3. *It is possible to solve n-consensus using only $\lceil n/\ell \rceil$ ℓ-buffers.*

6.2 A Lower Bound

In this section, we show a lower bound on the number of memory locations necessary for solving obstruction-free binary consensus among $n \geq 2$ processes.

A location r is *covered* by a process p in some configuration, if p is poised to perform ℓ-*buffer-write* on r. An ℓ-buffer is k-*covered* by a set of processes $\mathcal{P}$ in some configuration, if there are exactly k processes in $\mathcal{P}$ that cover it. A configuration is *at most k-covered* by $\mathcal{P}$, if no ℓ-buffer in the system is k'-covered by $\mathcal{P}$ in this configuration, for any $k' > k$.

Let $\mathcal{Q}$ be a set of processes, each of which is poised to perform ℓ-*buffer-write* in some configuration C. A *block write* by $\mathcal{Q}$ *from* C is an execution, starting at C, in which each process in $\mathcal{Q}$ takes exactly one step. If a block write is performed that includes ℓ different ℓ-*buffer-write* instructions to some ℓ-buffer, then any process that performs ℓ-*buffer-read* on that ℓ-buffer immediately afterwards, gets the same result (and ends up in the same state) regardless of the value of that ℓ-buffer in C.

We say that a set of processes $\mathcal{P}$ *can decide* $v \in \{0,1\}$ *from* a configuration C if there exists a $\mathcal{P}$-only execution from C in which v is decided. If $\mathcal{P}$ can decide both 0 and 1 from C, then $\mathcal{P}$ is *bivalent* from C. If $\mathcal{P}$ can decide v, but not $\bar{v}$, from C, then $\mathcal{P}$ is v-univalent from C. Note that, because the system is deterministic, each process p has only one solo execution starting from C that decides a value. Thus $\{p\}$ is either 0-univalent or 1-univalent from C.

To obtain the lower bound, we extend the proof of the $n-1$ lower bound on the number of registers required for solving n-process consensus [Zhu16], borrowing intuition about reserving executions from the $\Omega(n)$ lower bound for anonymous consensus [Gel15]. The following auxiliary lemmas are largely unchanged from [Zhu16]. The main difference is that we only perform block writes on ℓ-buffers that are ℓ-covered by $\mathcal{P}$.

Lemma 6.4. *There is an initial configuration from which the set of all processes in the system is bivalent.*

Proof. Consider an initial configuration I with two processes p_v for $v \in \{0,1\}$, where p_v starts with input v. Observe that $\{p_v\}$ is v-univalent from I since, initially, I is indistinguishable from the configuration where every process starts with v to p_v. Thus, $\{p_0, p_1\}$ is bivalent from I and, therefore, so is the set of all processes. □

Lemma 6.5. *Let R be a set of locations that are ℓ-covered by $\mathcal{P}$ in C, let $\mathcal{Q} \subseteq \mathcal{P}$ be the set of the $\ell \cdot |R|$ processes that cover R, and let β be a block write from C by $\mathcal{Q}$. Suppose some process $q \in \mathcal{Q}$ can decide some value $v \in \{0,1\}$ from configuration $C\beta$ and γ is a p-only execution from C by a process $p \notin \mathcal{Q}$ in which $\bar{v}$ is decided. Let γ' be the longest prefix of γ such that q can decide v from $C\gamma'\beta$. Then, in configuration in $C\gamma'$, p covers a location not in R and $\{p,q\}$ is bivalent from $C\gamma'\beta$.*

Proof. Let δ be the next step by p in γ after γ'. If δ is an *ℓ-buffer-write* to a location in R or is an *ℓ-buffer-read*, then $C\gamma'\beta$ and $C\gamma'\delta\beta$ are indistinguishable to q. Since q can decide v from $C\gamma'\beta$, but q is $\bar{v}$-univalent from $C\gamma'\delta\beta$, δ must be an *ℓ-buffer-write* to a location not in R. Thus $C\gamma'\beta\delta$ is indistinguishable from $C\gamma'\delta\beta$ to process q, so q is $\bar{v}$-univalent from $C\gamma'\beta\delta$. Hence, $\{p,q\}$ can decide $\bar{v}$ from $C\gamma'\beta$. By definition of γ', q can decide v from $C\gamma'\beta$. Thus, $\{p,q\}$ is bivalent from $C\gamma'\beta$, as desired. □

The next result says that if a set of processes is bivalent in some configuration, then it is possible to reach a configuration from which some process can decide 0 and another process can decide 1. It does not depend on what instructions are supported by the memory.

Lemma 6.6. *Suppose a set U of at least two processes is bivalent from configuration C. Then it is possible to reach, via a U-only execution from C, a configuration C' in which there are processes $p_0, p_1 \in U$, such that for each $v \in \{0,1\}$, p_v can decide v from C'.*

Proof. Let $\mathcal{V}$ be the set of all configurations from which U is bivalent and which are reachable from C by a U-only execution. Let k be the smallest integer such that there exist a configuration $C' \in \mathcal{V}$ and a set of processes $U' \subseteq U$ that is bivalent from C' with $|U'| = k$. Pick any such $C' \in \mathcal{V}$ and the corresponding set of processes U'.

Since a set containing only one process is either 0-univalent or 1-univalent from C', it follows that $|U'| \geq 2$. Consider a process $p \in U'$ and let $\overline{P} = U' \setminus \{p\}$ be the set of remaining processes in U'. Since $|\overline{P}| = k-1$, $\overline{P}$ must be v-univalent from C' for some $v \in \{0,1\}$. It follows that each process $q \in \overline{P}$ can decide v from C'.

If p can decide $\bar{v}$ from C' we are done. So, suppose that, like $\overline{P}$, p is v-univalent from C'.

Since U' is bivalent from C', there is a U'-only execution α from C' that decides $\bar{v}$. Let α' be the longest prefix of α such that both p and $\overline{P}$ are both v-univalent from $C'\alpha'$. Note that $\alpha' \neq \alpha$, because $\bar{v}$ is decided in α. Let δ be the next step in α after α'. Then either p or $\overline{P}$ can decide $\bar{v}$ from $C'\alpha'\delta$.

If δ is a step by a process in $\overline{P}$, then, since $\overline{P}$ is v-univalent from $C'\alpha'$, $\overline{P}$ is also v-univalent from $C'\alpha'\delta$. Therefore, p can decide $\bar{v}$ from $C'\alpha'\delta$. We are done because each process $q \in \overline{P}$ can decide v from $C'\alpha'\delta$.

Finally, suppose that δ is a step by p. Then, since p is v-univalent from $C'\alpha'$, p is also v-univalent from $C'\alpha'\delta$. Therefore, $\overline{P}$ can decide $\bar{v}$ from $C'\alpha'\delta$. However, $|\overline{P}| = k-1$. By definition of k, $\overline{P}$ is not bivalent from $C'\alpha'\delta$. Therefore $\overline{P}$ is $\bar{v}$-univalent from $C'\alpha'\delta$. This implies that every process $q \in \overline{P}$ can decide $\bar{v}$ from $C'\alpha'\delta$. Since we know $p \neq q$ can decide v from this configuration, the proof is complete. □

An induction similar to [Zhu16] allows the processes to reach a configuration that is at most ℓ-covered by a set of processes $\mathcal{R}$, while another process $z \notin \mathcal{R}$ covers a location that is not ℓ-covered by $\mathcal{R}$. This implies that the configuration is also at most ℓ-covered by $\mathcal{R} \cup \{z\}$, allowing the inductive step to go through.

Lemma 6.7. *Let C be a configuration and let $\mathcal{P}$ be a set of $n \geq 2$ processes. If $\mathcal{P}$ is bivalent from C, then there is a $\mathcal{P}$-only execution α and a pair of processes $\mathcal{Q} \subseteq \mathcal{P}$ such that $\mathcal{Q}$ is bivalent from $C\alpha$ and $C\alpha$ is at most ℓ-covered by the remaining processes $\mathcal{P} - \mathcal{Q}$.*

Proof. By induction on $|\mathcal{P}|$. The base case is when $|\mathcal{P}| = 2$, and holds with the empty execution. Now suppose $|\mathcal{P}| > 2$ and the claim holds for $|\mathcal{P}| - 1$. By Lemma 6.6, there is a $\mathcal{P}$-only execution γ a pair of processes $\{p_0, p_1\} \subset \mathcal{P}$ that are bivalent from $D = C\gamma$. Pick any process $z \in \mathcal{P} - \{p_0, p_1\}$. Then $\mathcal{P} - \{z\}$ is bivalent from D.

We construct a sequence of configurations $D_0, D_1, \ldots$ reachable from D, such that, for all $i \geq 0$, the following properties hold:

1. there exists a pair of processes $\mathcal{Q}_i \subseteq \mathcal{P} - \{z\}$ such that $\mathcal{Q}$ is bivalent from D_i,
2. D_i is at most ℓ-covered by the remaining processes $\mathcal{R}_i = (\mathcal{P} - \{z\}) - \mathcal{Q}_i$,
3. D_{i+1} is reachable from D_i by a $(\mathcal{P} - \{z\})$-only execution α_i which contains a block write β_i to the locations in D_i which are ℓ-covered by processes in $\mathcal{R}_i$.

We can construct D_0 by applying the induction hypothesis on D and $\mathcal{P} - \{z\}$, so that we obtain a $(\mathcal{P} - \{z\})$-only execution η such that the aforementioned properties hold in $D_0 = D\eta$. Now suppose we have D_i, $\mathcal{Q}_i$, $\mathcal{R}_i$ as defined previously. By Lemma 6.5 there is a $\mathcal{Q}_i$-only execution φ_i such that $\mathcal{P} - \{z\}$ is bivalent from $D_i\varphi_i\beta_i$, where β_i is a block write to the locations in D_i which are ℓ-covered by processes in $\mathcal{R}_i$. Applying the induction hypothesis on $D_i\varphi_i\beta_i$ and $\mathcal{P} - \{z\}$, there is a $(\mathcal{P} - \{z\})$-only execution ψ_i such that $D_{i+1} = D_i\varphi_i\beta_i\psi_i$ satisfies the aformentioned properties. Let $\alpha_i = \varphi_i\beta_i\psi_i$.

Since there are only finitely many locations, there exists $0 \leq i < j$ such that $\mathcal{R}_i$ covers the same set of registers in D_i as $\mathcal{R}_j$ does in D_j. We now insert steps of z so that no process in $\mathcal{P} - \{z\}$ can detect them. Consider any $\{z\}$-only execution ζ from $D_i\varphi_i$ that decides a value $v \in \{0, 1\}$. Since $D_i\varphi_i\beta_i$ is bivalent for $\mathcal{P} - \{z\}$, and β_i writes to ℓ-covered locations, ζ must contain a *ℓ-buffer-write* to a location that is not ℓ-covered by $\mathcal{R}_i$. Otherwise, $D_i\varphi_i\zeta\beta_i$ is indistinguishable from $D_i\varphi_i\beta_i$ to processes in $\mathcal{P} - \{z\}$, hence they can decide $\overline{v}$ from $D_i\zeta\beta_i$, which is impossible. Let ζ' be the longest prefix of ζ containing only writes to locations ℓ-covered by $\mathcal{R}_i$ in D_i. It follows that, in $D_i\varphi_i\zeta'$, z is poised to perform an *ℓ-buffer-write* to a location not ℓ-covered by $\mathcal{R}_i$ in D_i and, hence, $\mathcal{R}_j$ in D_j.

$D_i\varphi_i\zeta'\beta_i$ is indistinguishable from $D_i\varphi_i\beta_i$ to $\mathcal{P} - \{z\}$, so the $(\mathcal{P} - \{z\})$-only execution $\psi_i\alpha_{i+1}\cdots\alpha_{j-1}$ is applicable at $D_i\varphi_i\zeta'\beta_i$. Let $\alpha = \gamma\eta\alpha_0\cdots\alpha_{i-1}\varphi_i\zeta'\beta_i\psi_i\alpha_{i+1}\cdots\alpha_{j-1}$. Every process in $\mathcal{P} - \{z\}$ is in the same state in $C\alpha$ as it is in D_j. In particular, $\mathcal{Q}_j \subseteq \mathcal{P} - \{z\}$ is bivalent from D_j and, hence, from $C\alpha$, and every location in D_j is at most ℓ-covered by $\mathcal{R}_j = (P - \{z\}) - \mathcal{Q}_j$ in D_j and, hence, $C\alpha$. Moreover, since z takes no steps after $D_i\varphi\zeta'$, in $C\alpha$, z covers a register not ℓ-covered by $\mathcal{R}_i$ in D_i and, hence, $\mathcal{R}_j$ in D_j or $C\alpha$. Therefore, every location is at most ℓ-covered by $\mathcal{R}_j \cup \{z\} = \mathcal{P} - \mathcal{Q}_j$ in $C\alpha$. □

Finally, we can prove the main theorem.

Theorem 6.8. *Consider a memory consisting of ℓ-buffers. Then any algorithm for solving obstruction-free binary consensus for n processes uses at least $\lceil (n-1)/\ell \rceil$ locations.*

Proof. Consider an obstruction-free binary consensus algorithm. Let C be an initial configuration from which the set of all n processes, $\mathcal{P}$, is bivalent. Such a configuration exists by Lemma 6.4. Lemma 6.7 implies there is a reachable configuration C and a pair of processes $\mathcal{Q} \subseteq \mathcal{P}$ that are bivalent from C. Furthermore, C is at most ℓ-covered by the remaining processes $\mathcal{R}$. By the pigeonhole principle, $\mathcal{R}$ covers at least $\lceil (n-2)/\ell \rceil$ different locations. If $\lceil (n-2)/\ell \rceil < \lceil (n-1)/\ell \rceil$, then $n-2$ is a multiple of ℓ and every location covered by $\mathcal{R}$ is in fact ℓ-covered by $\mathcal{R}$. By Lemma 6.5, since $\mathcal{Q}$ is bivalent from C, we can use a process in $\mathcal{Q}$ to cover a location not ℓ-covered by $\mathcal{R}$. Hence, there are at least $\lceil (n-2)/\ell \rceil + 1 = \lceil (n-1)/\ell \rceil$ locations. □

The lower bound can be easily generalized to consensus algorithms that satisfy nondeterministic solo termination [FHS98]. It can also be extended to a *heterogeneous setting*, where memory locations need not be identical, for example, ℓ-buffers for possibly different values of ℓ. For this, we need to extend the definition of at most ℓ-covered to, instead, require each ℓ-buffer to be covered by at most ℓ processes, for all values of ℓ. Then we consider block writes to a set of locations containing ℓ different *ℓ-buffer-write* operations to each

ℓ-buffer in the set. The general result is that, for any algorithm which solves consensus for n processes and satisfies non-deterministic solo termination, the sum of capacities of all buffers must be at least $n-1$.

The lower bound also applies to systems in which the output of every non-trivial instruction on a memory location does not depend on the value of that location and the output of any trivial instruction is a function of the sequence of the preceding ℓ non-trivial instructions performed on the location. This is because such instructions can be implemented by ℓ-*buffer-read* and ℓ-*buffer-write* instructions. We record each invocation of a non-trivial instruction using ℓ-*buffer-write*. The return values of these instructions can be determined without even looking at the memory. To implement a trivial instruction, we perform ℓ-*buffer-read*, which returns a sequence containing the description of the last ℓ non-trivial instructions performed on the location, which is sufficient to determine the correct output.

7 Multiple Assignment

With m-register multiple assignment, we can atomically write to m locations. This plays an important role in [Her91], as m-register multiple assignment can used to solve wait-free consensus for $2m-2$ processes, but not for $2m-1$ processes.

In this section, we explore whether multiple assignment could improve the space complexity of solving obstruction-free consensus. A practical motivation to this question is that obstruction-free multiple assignment can be easily implemented using s simple transaction.

We prove a lower bound that is similar to the lower bound in Section 6.2. Suppose ℓ-*buffer-read*() and ℓ-*buffer-write*(x) instructions are supported on every memory location in a system and, for any subset of locations, we are allowed to atomically perform one ℓ-*buffer-write* instruction per location. Then $\lceil n/2\ell \rceil$ locations are necessary for n processes to solve obstruction-free consensus. As in Section 6.2, this result can be generalized to nondeterministic solo termination, different sets of instructions, and heterogeneous settings.

The main difficulty in this lower bound is proving an analogue of Lemma 6.5. In the absence of multiple assignment, if β is a block ℓ-*buffer-write* to a set of ℓ-covered locations R and δ is an ℓ-*buffer-write* to a location not in R, then β and δ trivially commute (in the sense that the resulting configurations are indistinguishable to all processes). However, a multiple assignment δ may now atomically ℓ-*buffer-write* to many locations, including some in R. Thus, it is now possible for processes to distinguish between $\beta\delta$ and $\delta\beta$. Using a careful combinatorial argument, we show how to perform a block of multiple assignments $\beta^1\beta^2$, such that, in β^i, ℓ-*buffer-write* is invoked exactly ℓ times for each location in R, and is never invoked for any location not in R. Given this, we can show that $\beta^1\delta\beta^2$ and $\delta\beta^1\beta^2$ are indistinguishable to all processes, which is enough to prove an analogue of Lemma 6.5.

First, we define the notion of covering in this setting. We say that a process p *covers* location r in a configuration C, when the next step of p in C is a multiple assignment that involves an ℓ-*buffer-write* to r. The next definition is key to our proof. *A k-packing* of a set of processes $\mathcal{P}$ in some configuration C is a function π mapping each process in $\mathcal{P}$ to some memory location it covers such that no location r has more than k processes mapped to it (i.e., $|\pi^{-1}(r)| \leq k$). When $\pi(p) = r$ we say that π *packs* p in r. A k-packing may not always exist, or there may be many, depending on the configuration, the set of processes, and the value of k. A location r is *fully k-packed* by $\mathcal{P}$ in configuration C, if a k-packing of $\mathcal{P}$ exists in C and all such k-packings pack exactly k processes in r.

Suppose that, in some configuration, there are two k-packings of the same set of processes such that the first packs more processes in some location r than the second. We show there is a location r' in which the first packing packs fewer processes than the second and a k-packing that pack one less process to location r, one more process to location r' and the same number of processes to all other locations, as compare to the first packing. This iterative procedure has some resemblance to cuckoo hashing [PR01]. The proof relies on existence of a certain Eulerian path in a multigraph that we build to represent these two k-packings.

Lemma 7.1. *Suppose g and h are two k-packings of the same set of processes $\mathcal{P}$ in some configuration C and r_1 is a location such that g packs more processes in r_1 than h does (i.e., $|g^{-1}(r_1)| > |h^{-1}(r_1)|$). Then, there exists a sequence of locations $r_1, r_2, \ldots, r_t$ and a sequence of distinct processes $p_1, p_2, \ldots, p_{t-1}$, such that h packs more processes in r_t than g, (i.e., $|h^{-1}(r_t)| > |g^{-1}(r_t)|$) and, for $1 \leq i \leq t-1$, $g(p_i) = r_i$ and $h(p_i) = r_{i+1}$.*

Proof. Consider a multigraph with one node for each memory location in the system and one directed edge labelled by p, from node $g(p)$ to node $h(p)$, for each process $p \in \mathcal{P}$. Therefore, the in-degree of a node v is equal to $|h^{-1}(v)|$, which is the number of processes that are packed into memory location v by h, and the out-degree of node v is equal to $|g^{-1}(v)|$, which is the number of processes that are packed in v by g.

Now, consider any maximal Eulerian path in this multigraph starting from node r_1. This path follows a sequence of edges for as long as possible without repeating any edge. Let $r_1, \ldots, r_t$ be the sequence of nodes visited (which may contain repetitions) and let p_i be the labels of the traversed edges, in order. Then $g(p_i) = r_i$ and $h(p_i) = r_{i+1}$ for $1 \leq i \leq t-1$. Since the path is Eulerian, the labels of the edges are distinct, which, in turn, guarantees that these sequences are finite. Finally, by maximality, the last node in the sequence must have more incoming than outgoing edges, so $|h^{-1}(r_t)| > |g^{-1}(r_t)|$. □

Corollary 7.2. *Let the k-packings g and h and the sequences r_i and p_i be defined as in Lemma 7.1. For $1 \leq j < t$, there exists a k-packing g', such that g' packs one less process than g in r_j, one more process than g in r_t, and the same number of processes as g in all other locations.*

Proof. We construct g' from g by re-packing each process p_i from r_i to r_{i+1} for all $j \leq i < t$. Then $g'(p_i) = r_{i+1}$ for $j \leq i < t$ and $g'(p) = g(p)$ for all other processes p. Notice that p_i covers r_{i+1}, since $h(p_i) = r_{i+1}$ and h is a k-packing.

As compared to g, g' packs one less process in r_j, one more process in r_t, and the same number of processes in every other location. Since h is a k-packing, it packs at most k processes in r_t. Because g is a k-packing that packs less processes in r_t, g' is also a k-packing. □

Let $\mathcal{Q}$ be a set of processes, each of which is poised to perform a multiple assignment in some configuration C. A *block multi-assignment* by $\mathcal{Q}$ from C is an execution interval starting at C, in which each process in $\mathcal{Q}$ takes exactly one step. Consider some configuration C, a set of processes $\mathcal{P}$, and a 2ℓ-packing π of $\mathcal{P}$ in C. Let R be the set of all locations that are fully 2ℓ-packed by $\mathcal{P}$ in C. By definition, π packs exactly 2ℓ processes from $\mathcal{P}$ in each location $r \in R$.

Suppose $\mathcal{Q}_1, \mathcal{Q}_2 \subseteq \mathcal{P}$ is a partition of all processes that are packed in R by π, such that for each location $r \in R$, exactly ℓ packed processes belong to $\mathcal{Q}_1$ and the remaining ℓ packed processes belong to $\mathcal{Q}_2$. Let β_i be a block multi-assignment by $\mathcal{Q}_i$ for $i \in \{1, 2\}$.

Notice that, for any location $r \in R$, after β_1 or β_2, the outcome of any subsequent *ℓ-buffer-read* on r does not depend on multiple assignments that occurred prior to the block multi-assignment. Moreover, we can prove the following crucial property about these block multi-assignments to fully packed locations.

Lemma 7.3. *Neither β_1 nor β_2 involves an ℓ-buffer-write to a location outside of R.*

Proof. Assume the contrary. Let $q \in \mathcal{Q}_1 \cup \mathcal{Q}_2$ be a process with $\pi(q) \in R$, such that q also covers some location $r_1 \notin R$ in C. It must be the case that π packs exactly 2ℓ processes in r_1 ($|\pi^{-1}(r_1)| = 2\ell$). Otherwise, re-packing q in r_1 instead of $\pi(q)$ leads to another 2ℓ-packing that does not pack exactly 2ℓ processes in $\pi(q) \in R$, contradicting the definition of a fully 2ℓ-packed location.

Since R is the set of all fully 2ℓ-packed locations, there exists a 2ℓ-packing h, which packs strictly less than 2ℓ processes in r_1. Applying Lemma 7.1 for 2ℓ-packings π and h, we get a sequence of locations $r_1, \ldots, r_t$ and a sequence of processes $p_1, \ldots p_{t-1}$. $|\pi^{-1}(r_t)| < |h^{-1}(r_t)|$, and since h is a 2ℓ-packing, we must have that $|\pi^{-1}(r_t)| < 2\ell$, i.e. π packs strictly less than 2ℓ processes in r_t in C. We consider two cases.

First, assume that q does not occur in the sequence, i.e. $\forall i : p_i \neq q$. We know that $|\pi^{-1}(r_1)| = 2\ell$ and $|\pi^{-1}(r_t)| < 2\ell$, implying that the locations r_1 and r_t must be different. This allows us to apply Corollary 7.2 for the whole sequence with $j = 1$. We get a 2ℓ-packing π' that packs less than 2ℓ processes in r_1, and exactly 2ℓ processes in each of the fully 2ℓ-packed locations R, including $\pi(q)$. Moreover, we did not re-pack process q, so $\pi'(q) = \pi(q)$. Hence, modifying π' by re-packing q in r_1 instead of $\pi'(q)$ again leads to a 2ℓ-packing that does not pack exactly 2ℓ processes in a location in R, a contradiction.

Now, assume that $q = p_s$, for some s. By the properties of our sequences, we know $r_s = \pi(p_s) = \pi(q) \in R$, and since this location is fully 2ℓ-packed in configuration C we have $|\pi^{-1}(r_s)| = 2\ell$. Similar to the first case, since $|\pi^{-1}(r_t)| < 2\ell$ in C, locations r_s and r_t must be different and we now apply Corollary 7.2 with $j = s$. We get a 2ℓ-packing π'' that packs less than 2ℓ processes in location $r_s \in R$. This is a contradiction with the definition of a fully 2ℓ-packed location, completing the proof. □

We can now prove a lemma that replaces Lemma 6.5 in the main argument.

Lemma 7.4. *Suppose $C\beta^1\beta^2$ is v-deciding for $\mathcal{Q}_1 \cup \mathcal{Q}_2$ for some $v \in \{0,1\}$. If γ is a solo execution from C by a process $p \notin \mathcal{Q}_1 \cup \mathcal{Q}_2$ in which $\overline{v}$ is decided, and γ' is the longest prefix of γ such that $C\gamma'\beta^1\beta^2$ is v-deciding for $\mathcal{Q}_1 \cup \mathcal{Q}_2$, then p is poised to perform a multiple assignment which contains an ℓ-buffer-write to a location not in R in $C\gamma'$ and the configuration $C\gamma'\beta^1$ is bivalent for $\mathcal{Q}_1 \cup \mathcal{Q}_2 \cup \{p\}$.*

Proof. Let δ be the next step by p in γ after γ'. Since $C\gamma'\beta^1\beta^2$ is v-deciding for $\mathcal{Q}_1 \cup \mathcal{Q}_2$, but $C\gamma'\delta\beta^1\beta^2$ is $\bar{v}$-univalent for $\mathcal{Q}_1 \cup \mathcal{Q}_2$, δ must be a multiple assignment which includes an *ℓ-buffer-write* to a location not in R. (If δ is a read, the resulting configurations would be indistinguishable. Similarly, if δ is a multiple assignment containing only *ℓ-buffer-write*'s to locations in R, then no process in $\mathcal{Q}_1 \cup \mathcal{Q}_2$ can observe a difference due to the block writes $\beta^1\beta^2$.)

We claim that the configuration $C\gamma'\beta^1\delta\beta^2$ is indistinguishable from $C\gamma'\delta\beta^1\beta^2$ to $\mathcal{Q}_1 \cup \mathcal{Q}_2$. Indeed, for each location $r \in R$, since $|\pi^{-1}(r) \cap \mathcal{Q}_i| = \ell$ (that is, r is covered by exactly ℓ processes in $\mathcal{Q}_i$), the contents in location r is the same in $C\gamma'\delta\beta^1\beta^2$ as it is in $C\gamma'\beta^1\delta\beta^2$. On the other hand, for each location $r \notin R$, by Lemma 7.3, neither β^1 nor β^2 performs an *ℓ-buffer-write* to r, thus, the contents in location r is the same in $C\gamma'\delta\beta^1\beta^2$ as it is in $C\gamma'\beta^1\delta\beta^2$.

Therefore, $C\gamma'\beta^1\delta\beta^2$ is $\bar{v}$-univalent for $\mathcal{Q}_1 \cup \mathcal{Q}_2$, hence $C\gamma'\beta^1$ is $\bar{v}$-deciding for $\mathcal{Q}_1 \cup \mathcal{Q}_2 \cup \{p\}$. Moreover, $C\gamma'\beta^1\beta^2$ is v-deciding for $\mathcal{Q}_1 \cup \mathcal{Q}_2$ by definition of γ', thus $C\gamma'\beta^1$ is v-deciding for $\mathcal{Q}_1 \cup \mathcal{Q}_2 \cup \{p\}$. We have thus established the desired bivalency of $C\gamma'\beta^1$, completing the proof. □

Using these tools, we can prove the following analogue of Lemma 6.7:

Lemma 7.5. *Let C be a configuration and let $\mathcal{P}$ be a set of $n \geq 2$ processes. If $\mathcal{P}$ is bivalent from C, then there is a $\mathcal{P}$-only execution α and a pair of processes $\mathcal{Q} \subseteq \mathcal{P}$ such that $\mathcal{Q}$ is bivalent from $C\alpha$ and there is a 2ℓ-packing π of the remaining processes $\mathcal{P} - \mathcal{Q}$ in $C\alpha$.*

Proof. By induction on $|\mathcal{P}|$. The base case is when $|\mathcal{P}| = 2$, and holds with the empty execution. Now suppose $|\mathcal{P}| > 2$ and the claim holds for $|\mathcal{P}| - 1$. By Lemma 6.6, there is a $\mathcal{P}$-only execution γ a pair of processes $\{p_0, p_1\} \subset \mathcal{P}$ that are bivalent from $D = C\gamma$. Pick any process $z \in \mathcal{P} - \{p_0, p_1\}$. Then $\mathcal{P} - \{z\}$ is bivalent from D.

We construct a sequence of configurations $D_0, D_1, \ldots$ reachable from D, such that, for all $i \geq 0$, the following properties hold:

1. there exists a pair of processes $\mathcal{Q}_i \subseteq \mathcal{P} - \{z\}$ such that $\mathcal{Q}$ is bivalent from D_i,
2. there is a 2ℓ-packing π_i of the remaining processes $\mathcal{R}_i = (\mathcal{P} - \{z\}) - \mathcal{Q}_i$ in D_i,
3. D_{i+1} is reachable from D_i by a $(\mathcal{P} - \{z\})$-only execution α_i, and
4. if R_i is the set of all fully 2ℓ-packed locations by $\mathcal{R}_i$ in D_i, then α_i contains a block multi-assignment β_i such that, for each location $r \in R_i$, β_i contains at least ℓ multiple assignments which perform *ℓ-buffer-write* to r.

We can construct D_0 by applying the induction hypothesis on D and $\mathcal{P} - \{z\}$, so that we obtain a $(\mathcal{P} - \{z\})$-only execution η such that the aforementioned properties hold in $D_0 = D\eta$. Now suppose we have D_i, $\mathcal{Q}_i$, $\mathcal{R}_i$, π_i, and R_i as defined previously. By Lemma 7.4, if we let $\mathcal{R}_i^1$ be the set of all the first ℓ processes that are packed into each location in R_i in D_i by π_i, and $\mathcal{R}_i^2$ be the remaining processes that are packed into R_i in D_i, then there is a $\mathcal{Q}_i$-only execution φ such that $D_i\varphi_i\beta_i$ is bivalent for $\mathcal{R}_i^1 \cup \mathcal{R}_i^2 \cup \mathcal{Q}_i \subseteq \mathcal{P} - \{z\}$, where β_i is a block multi-assignment by $\mathcal{R}_i^1$. By definition of $\mathcal{R}_i^1$, for each location $r \in R_i$, β_i contains at least ℓ multiple assignments which perform *ℓ-buffer-write* to r. Applying the induction hypothesis on $D_i\varphi_i\beta_i$ and $\mathcal{P} - \{z\}$, there is a $(\mathcal{P} - \{z\})$-only execution ψ_i such that $D_{i+1} = D_i\varphi_i\beta_i\psi_i$ satisfies the aforementioned properties. Let $\alpha_i = \varphi_i\beta_i\psi_i$.

Since there are only finitely many locations, there exists $0 \leq i < j$ such that $R_i = R_j$. That is, the set of fully 2ℓ-packed locations by $\mathcal{R}_i$ in D_i is the same as the set of fully 2ℓ-packed locations by $\mathcal{R}_j$ in D_j. Let $R = R_i$. We now insert steps of z so that no process in $\mathcal{P} - \{z\}$ can detect them. Consider any $\{z\}$-only execution ζ from $D_i\varphi_i$ that decides a value $v \in \{0,1\}$. Since $D_i\varphi_i\beta_i$ is bivalent for $\mathcal{P} - \{z\}$, and, for each location $r \in R$, β_i contains at least ℓ multiple assignments which perform *ℓ-buffer-write* to r, ζ must contain

a ℓ-*buffer-write* to a location not in R. Otherwise, $D_i\varphi_i\zeta\beta_i$ is indistinguishable from $D_i\varphi_i\beta_i$ to processes in $\mathcal{P}-\{z\}$, hence they can decide $\overline{v}$ from $D_i\zeta\beta_i$, which is impossible. Let ζ' be the longest prefix of ζ containing only writes to locations in R. It follows that, in $D_i\varphi_i\zeta'$, z is poised to perform an ℓ-*buffer-write* to a location not in R in D_i and, hence, D_j.

$D_i\varphi_i\zeta'\beta_i$ is indistinguishable from $D_i\varphi_i\beta_i$ to $\mathcal{P}-\{z\}$, so the $(\mathcal{P}-\{z\})$-only execution $\psi_i\alpha_{i+1}\cdots\alpha_{j-1}$ is applicable at $D_i\varphi_i\zeta'\beta_i$. Let $\alpha = \gamma\eta\alpha_0\cdots\alpha_{i-1}\varphi_i\zeta'\beta_i\psi_i\alpha_{i+1}\cdots\alpha_{j-1}$. Every process in $\mathcal{P}-\{z\}$ is in the same state in $C\alpha$ as it is in D_j. In particular, $\mathcal{Q}_j \subseteq \mathcal{P}-\{z\}$ is bivalent from D_j and, hence, from $C\alpha$. Furthermore, the 2ℓ-packing π_j of $\mathcal{R}_j$ in D_j is a 2ℓ-packing of $\mathcal{R}_j$ in $C\alpha$, and the set of locations that are fully 2ℓ-packed by $\mathcal{R}_j$ in $C\alpha$ is R. Since z takes no steps after $D_i\varphi\zeta'$, in $C\alpha$, z covers a location r not in R. By definition of R, since $r \notin R$, there is a 2ℓ-packing π'_j of $\mathcal{R}_j$ in $C\alpha$ which packs less than 2ℓ processes into r. Thus, we can define a 2ℓ-packing π of $\mathcal{R}_j \cup \{z\}$ by packing each process in $\mathcal{R}_j$ according to π'_j and packing z into r. It follows that π is a 2ℓ-packing of $\mathcal{P}-\mathcal{Q}_j = \mathcal{R}_j \cup \{z\}$ in $C\alpha$. □

We can now prove the main theorem.

Theorem 7.6. *Consider a memory consisting of ℓ-buffers. If the processes can atomically ℓ-buffer-write to any subset of the ℓ-buffers, then any algorithm for solving obstruction consensus for n processes needs at least $\lceil (n-1)/2\ell \rceil$ space.*

Proof. Consider an obstruction-free binary consensus algorithm. Let C be an initial configuration from which the set of all n processes, $\mathcal{P}$, is bivalent. Such a configuration exists by Lemma 6.4. Lemma 7.5 implies there is a reachable configuration C and a pair of processes $\mathcal{Q} \subseteq \mathcal{P}$ that are bivalent from C. Furthermore, there is a 2ℓ-packing π of the remaining processes $\mathcal{R} = \mathcal{P}-\mathcal{Q}$ in C. Let R be the set of all locations that are covered by $\mathcal{R}$. By the pigeonhole principle, $|R| \geq \lceil (n-2)/2\ell \rceil$ different locations. If $|R| < \lceil (n-1)/2\ell \rceil$, then $|R| = \lceil (n-2)/2\ell \rceil$, $n-2$ is a multiple of 2ℓ, and every location in R is fully 2ℓ-packed by $\mathcal{R}$. By Lemma 7.4, since $\mathcal{Q}$ is bivalent from C, we can use a process in $\mathcal{Q}$ to cover a location not fully 2ℓ-packed by $\mathcal{R}$. Hence, there are at least $\lceil (n-2)/2\ell \rceil + 1 = \lceil (n-1)/2\ell \rceil$ locations. □

8 Swap and Read

It is possible to solve n-consensus using $n-1$ memory locations that support only *read*() and *swap*(x) The idea is to modify the anonymous algorithm of [Zhu15], which uses n locations supporting *read*() and *write*(x). In that algorithm, processes obtain a snapshot of all n values by repeatedly performing collect. If all n memory locations have the same (non-$\perp$) contents, the process returns the value each of them contains. Otherwise, it overwrites one location and repeats the same sequence of steps again. After some process p returns, each of the remaining $n-1$ processes may overwrite one location before taking a snapshot. This ensures that that the value p decided is not lost, since it is contained in at least one location.

In the modified algorithm, instead of collecting the contents from n locations, a process collects the contents from $n-1$ locations and also uses the result of its last *swap*. After some process p returns, if all $n-1$ remaining processes perform *swap* on different locations, the result of each of their *swap* operations is the contents of each location seen by p in its last snapshot. Otherwise, as in the original algorithm, the value p decided is contained in at least one location.

A lower bound of $\Omega(\sqrt{n})$ locations when the system only supports *swap* and*read* instructions appears in [FHS98]. We can also adapt the recent proof of [Zhu16], which corresponds to our proof for buffered read-write instructions with capacity $\ell = 1$, to get a simpler proof of the $\Omega(\sqrt{n})$ lower bound

Only the proof of Lemma 6.7 needs to change. It must change because the processes perform block swaps instead of block writes, and can detect the steps by process z. To account for this, in the inductive step from i to $i+1$ locations, we consider a set Z of $i+1$ processes instead of the single process z. We still perform the block swaps and reach new configurations. If, in any such configuration, two processes in Z cover the same location. then steps are allocated to one of them, until it covers a location not covered by any other process in Z. This is repeated until all processes in Z cover $i+1$ different locations. Counting the number of extra processes used in the argument leads to the lower bound of $\Omega(\sqrt{n})$ locations.

9 Test-and-Set and Read

Consider a system that supports only *test-and-set*() and *read*(). If there are only 2 processes, then it is possible to solve wait-free consensus using a single memory location. However, we claim that any algorithm for solving obstruction-free binary consensus among $n \geq 3$ processes must use an unbounded number of memory locations. The key is to prove the following analogue of Lemma 6.5.

Lemma 9.1. *Let P be a set of at least 3 processes and let C be a configuration. If P is bivalent from C, then, for every $k \geq 0$, there exists a P-only execution α_k from C such that P is bivalent from $C\alpha_k$ and at least k locations have been set to 1 in $C\alpha_k$.*

Proof. By induction on k. The base case is when $k = 0$ and is trivial. Given α_k, we show how to construct α_{k+1}. By Lemma 6.6, there is a P-only execution φ from $C\alpha_k$ and a pair of processes, p_0 and p_1, that can decide 0 and 1, respectively, from $C\alpha_k\varphi$. Pick $z \in P - \{p_0, p_1\}$. Since $\{p_0, p_1\}$ is bivalent from $C\alpha_k\varphi$, we may run z until it first performs a *test-and-set*() on a memory location, call it r, that had not been set to 1. Let β denote this execution. If $C\alpha_k\varphi\beta$ is bivalent from $\{p_0, p_1\}$, then we are done with $\alpha_{k+1} = \alpha_k\varphi\beta$. So, without loss of generality, suppose $\{p_0, p_1\}$ is 0-univalent from $C\alpha_k\varphi\beta$. Let ψ be the longest prefix of p_1's solo terminating execution from $C\alpha_k\varphi$ such that $C\alpha_k\varphi\psi\beta$ is 0-univalent for $\{p_0\}$. It follows that, in configuration $C\alpha_k\varphi\psi$, the set β' that p_1 is poised to perform is a *test-and-set*() on an unset memory location. Note that this location is not r, since, otherwise, $C\alpha_k\varphi\psi\beta'\beta$ is indistinguishable from $C\alpha_k\varphi\psi\beta$ to p_0. Therefore, $C\alpha_k\varphi\psi\beta\beta'$ is indistinguishable from $C\alpha_k\varphi\psi\beta'\beta$ to p_0. Let $\alpha_{k+1} = \alpha_k\varphi\psi\beta$. It follows that $\{p_0, p_1\}$ is bivalent from $C\alpha_{k+1}$. Furthermore, β includes a *test-and-set*() to a location that was not already set to 1 and $C\alpha_k$ had at least k locations set to 1, $C\alpha_{k+1}$ has at least $k + 1$ locations set to 1. □

By Lemma 6.4, there is an initial configuration from which the set of all processes in the system is bivalent. It follows that we may repeatedly apply the preceding lemma to show that any binary consensus algorithm for $n \geq 3$ processes uses an unbounded number of locations.

Theorem 9.2. *For $n \geq 3$, it is not possible to solve n-consensus using a bounded number of memory locations supporting only read() and test-and-set().*

There is an algorithm for obstruction-free binary consensus that uses an unbounded number of binary registers, initially all set to 0, and only performs *write*(1) [GR05]. The idea is to simulate a counter using an unbounded number of binary registers and then to run the racing counters algorithm presented in Lemma 3.1. In their algorithm, there are two unbounded tracks on which processes (with preference 0 and 1, respectively) race, flipping 0's to 1's one by one on their respective tracks to indicate progress. A process changes its preference if it sees that the number of 1's on its side is less than the number of 1's on the other side. Once a process sees that its side is at least 2 ahead of the other side, it decides its current preference.

It is possible to generalize this algorithm to solve n-valued consensus by having n race tracks, where register j is used for track k whenever $j \equiv k \bmod n$. Since *test-and-set*() is at least as powerful as *write*(1), we get the following result.

Theorem 9.3. *It is possible to solve n-consensus using only an unbounded number of memory locations supporting only read() and either write(1) or test-and-set().*

Now, suppose we can also *reset*() a memory location from 1 to 0 or perform *write*(0). Then we can solve n-consensus using only $O(n \log n)$ locations by combining an existing binary consensus algorithm that uses only $2n$ binary registers [Bow11] with Lemma 5.2. There is a slight subtlety in that we don't have *write*(x). However, in the proof of Lemma 5.2, it suffices to use n binary locations in each layer that are initialized to 0, instead of 2 registers. The idea is that a process sets the i'th location to 1 to indicate that it had value $i - 1$.

Theorem 9.4. *It is possible to solve n-consensus using only $O(n \log n)$ memory locations supporting only read(), write(1) or test-and-set(), and write(0) or reset().*

10 Conclusions and Future Work

In this paper, we defined a hierarchy based on the space complexity of solving consensus. We used consensus because it is a well-studied, general problem that seems to capture a fundamental difficulty of multiprocessor synchronization. Moreover, consensus is *universal*: any sequentially defined object can be implemented in a wait-free way using only consensus objects and registers [Her91].

We did not address the issue of universality within our hierarchy. One history object can be used to implement any sequentially defined object. Consequently, it may make sense to consider defining a hierarchy on sets of instructions based on their ability to implement a history object, a compare-and-swap object, or a repeated consensus object shared by n processes. However, the space complexity of solving n-consensus is the same as the space complexity of implementing these long-lived objects, for many of the instruction sets that we considered. (It appears that this may not be true for $\{read, increment\}$.)

It may be that a truly accurate complexity-based hierarchy would have to take time complexity into consideration. However it is even unclear which definition of time to use. Exploring this may be an important future direction.

There are several other interesting open problems. To the best of our knowledge, all existing space lower bounds rely on a combination of covering and indistinguishability arguments. However, when the covering processes apply $swap(x)$, as opposed to $write(x)$, they can observe differences between executions, so they can no longer be reused to maintain indistinguishability. This means that getting a tighter space lower bound for $\{swap(x), read()\}$ would most likely require a completely novel approach. An algorithm that uses less than $n-2$ space would be even more surprising, as the processes would necessarily have to adapt their access patterns to the memory locations based on the swapped values, in order to circumvent the argument from [Zhu16]. The authors are unaware of any such algorithm.

We conjecture that, for sets of instructions, $\mathcal{I}$, which contain only $read()$, $write(x)$, and either $increment()$ or $fetch\text{-}and\text{-}increment()$, $\mathcal{SP}(\mathcal{I}, n) \in \Theta(\log n)$. Similarly, we conjecture, for $\mathcal{I} = \{read(), write(0), write(1)\}$, $\mathcal{SP}(\mathcal{I}, n) \in \Theta(n \log n)$. Proving this is likely to require techniques that depend on the number of input values, such as in the lower bound for m-valued adopt-commit objects in [AE14].

We would like to understand the properties of sets of instructions at certain levels in the hierarchy. For instance, what properties enable a collection of instructions to solve n-consensus using a single location? Is there an interesting characterization of the sets of instructions $\mathcal{I}$ for which $\mathcal{SP}(\mathcal{I}, n)$ is constant?

How do subsets of a set of instructions relate to one another in terms of their locations in the hierarchy? Alternatively, what combinations of sets of instructions decrease the amount of space needed to solve consensus? For example, using only $read()$, $write(x)$, and either $increment()$ or $decrement()$, more than one memory location is needed to solve binary consensus. But with both $increment()$ and $decrement()$ a single location suffices. Are there general properties governing these relations?

11 Acknowledgments

Support is gratefully acknowledged from the Natural Science and Engineering Research Council of Canada, the National Science Foundation under grants CCF-1217921, CCF-1301926, and IIS-1447786, the Department of Energy under grant ER26116/DE-SC0008923, and Oracle and Intel corporations.

The authors would like to thank Michael Coulombe, Dan Alistarh, Yehuda Afek, Eli Gafni and Philipp Woelfel for helpful conversations and feedback.

References

[AAC09] James Aspnes, Hagit Attiya, and Keren Censor. Max registers, counters, and monotone circuits. In *Proc. 28th ACM Symposium on Principles of Distributed Computing*, pages 36–45, 2009.

[AAD+93] Yehuda Afek, Hagit Attiya, Danny Dolev, Eli Gafni, Michael Merritt, and Nir Shavit. Atomic snapshots of shared memory. *Journal of the ACM*, 40(4):873–890, 1993.

[AE14] James Aspnes and Faith Ellen. Tight bounds for adopt-commit objects. *Theory Comput. Syst.*, 55(3):451–474, 2014.

[AH90] James Aspnes and Maurice Herlihy. Fast randomized consensus using shared memory. *Journal of Algorithms*, 11(3):441–461, 1990.

[Bow11] Jack R. Bowman. Obstruction-free snapshot, obstruction-free consensus, and fetch-and-add modulo k. Technical Report TR2011-681, Computer Science Department, Dartmouth College, 2011. http://www.cs.dartmouth.edu/reports/TR2011-681.pdf.

[BRS15] Zohir Bouzid, Michel Raynal, and Pierre Sutra. Brief announcement: Anonymous obstruction-free (n, k)-set agreement with $n - k + 1$ atomic read/write registers. In *Proc. 29th International Symposium on Distributed Computing (DISC)*, page 669, 2015.

[FHS98] Faith Ellen Fich, Maurice Herlihy, and Nir Shavit. On the space complexity of randomized synchronization. *Journal of the ACM*, 45(5):843–862, 1998.

[FLMS05] Faith Ellen Fich, Victor Luchangco, Mark Moir, and Nir Shavit. Obstruction-free algorithms can be practically wait-free. In *Proc. 19th International Conference on Distributed Computing (DISC)*, pages 78–92, 2005.

[Gel15] Rati Gelashvili. On the optimal space complexity of consensus for anonymous processes. In *Proc. 29th International Symposium on Distributed Computing (DISC)*, pages 452–466, 2015.

[GHHW13] George Giakkoupis, Maryam Helmi, Lisa Higham, and Philipp Woelfel. An $\mathcal{O}(\sqrt{n})$ space bound for obstruction-free leader election. In *Proc. 27th International Symposium on Distributed Computing (DISC)*, pages 46–60, 2013.

[GR05] Rachid Guerraoui and Eric Ruppert. What can be implemented anonymously? In *Proc. 19th International Symposium on Distributed Computing (DISC)*, pages 244–259, 2005.

[Her91] Maurice Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 13(1):124–149, 1991.

[HR00] Maurice Herlihy and Eric Ruppert. On the existence of booster types. In *Proc. 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 653–663, 2000.

[HS12] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming*. Morgan Kaufmann, 2012.

[Int12] Intel. *Transactional Synchronization in Haswell*, 2012. http://software.intel.com/en-us/blogs/2012/02/07/transactional-synchronization-in-haswell.

[Jay93] Prasad Jayanti. On the robustness of herlihy's hierarchy. In *Proc. 12th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 145–157, 1993.

[LH00] Wai-Kau Lo and Vassos Hadzilacos. All of us are smarter than any of us: Nondeterministic wait-free hierarchies are not robust. *SIAM Journal on Computing*, 30(3):689–728, 2000.

[PR01] Rasmus Pagh and Flemming Friche Rodler. Cuckoo hashing. In *Proc. 9th Annual European Symposium on Algorithms (ESA)*, pages 121–133, 2001.

[Ray12] Michel Raynal. *Concurrent programming: algorithms, principles, and foundations*. Springer Science & Business Media, 2012.

[Rup00] Eric Ruppert. Determining consensus numbers. *SIAM Journal on Computing*, 30(4):1156–1168, 2000.

[Sch97] Eric Schenk. The consensus hierarchy is not robust. In *Proc. 16th ACM Symposium on Principles of Distributed Computing (PODC)*, page 279, 1997.

[Tau06] Gadi Taubenfeld. *Synchronization algorithms and concurrent programming*. Pearson Education, 2006.

[Zhu15] Leqi Zhu. Brief announcement: Tight space bounds for memoryless anonymous consensus. In *Proc. 29th International Symposium on Distributed Computing (DISC)*, page 665, 2015.

[Zhu16] Leqi Zhu. A tight space bound for consensus. In *Proc. 48th ACM Symposium on Theory of Computing (STOC)*, 2016. To appear.